Forecasting Hydrological Time Series Using a Temporal Fusion Transformer Model

Viktória Várkonyi, Zsolt Vizi
University of Szeged, Szeged, Hungary

Reliable short-term water level forecasting plays an important role in hydrology, especially in river sections where timely predictions can support flood protection, water management and operational decision-making. In this presentation, we focus on the implementation and evaluation of a Temporal Fusion Transformer model [1] applied to daily hydrological time series from the Tisza river system [2]. The aim of the study is to forecast the water level at the Szeged station over a multi-day prediction horizon, using past water-related observations together with future-known calendar information.

We first briefly introduce the forecasting problem and the structure of the dataset. The dataset contains daily stage and discharge measurements from several gauging stations. Training samples are created using a sliding-window approach, where each sample consists of a fixed-length historical input window followed by a prediction horizon. This formulation allows the model to produce multi-horizon forecasts instead of predicting only the next time step.

The main part of the talk presents the Python and PyTorch implementation. We describe how the key components of the Temporal Fusion Transformer were constructed in code, including the variable selection networks, the processing of static and time-dependent inputs, the LSTM-based temporal layers, and the interpretable multi-head attention mechanism. The model handles past observed variables, future-known variables and static information separately, which makes it suitable for structured hydrological forecasting tasks. The training procedure is based on quantile loss, allowing the model to provide not only a central forecast, but also lower and upper quantile estimates that represent prediction uncertainty.

Furthermore the presentation also focuses on the evaluation of the results. We show how the model is assessed across different forecast horizons using standard error metrics and visual comparisons between predicted and observed water levels. We also examine how the prediction error changes as the forecast horizon increases, how the errors are distributed, and how the model behaves across different water level ranges. Particular attention is given to higher water levels, since these are especially relevant from a flood-risk perspective. In addition, the TFT model is compared with an LSTM-based encoder-decoder baseline in order to discuss the strengths and limitations of the Transformer-based approach. Beyond numerical performance, we also briefly discuss the interpretability potential of the model, especially through its variable selection and attention mechanisms.

[1] B. Lim, S. Ö. Arik, N. Loeff, T. Pfister, Temporal fusion transformers for interpretable multi-horizon time series forecasting, *Int. J. Forecast.* **37** (2021), 1748–1764.

[2] Z. Vizi, B. Batki, L. Rátki, Sz. Szalánczi, I. Fehérváry, P. Kozák, T. Kiss, Water level prediction using long short-term memory neural network model for a lowland river: a case study on the Tisza River, Central Europe, *Environ. Sci. Eur.* **35** (2023), Article 92.